13th International Symposium on Turbulence and Shear Flow Phenomena (TSFP13)
Montreal, Canada, June 25–28, 2024

OPTIMISING WIND TURBINE PLACEMENT IN HILLY TERRAINS: INSIGHTS FROM BAYESIAN OPTIMISATION BASED ON LARGE EDDY SIMULATIONS

C. Jané-Ippel
Department of Aeronautics
Imperial College London
London SW7 2AZ, UK
christian.jane-ippel19@imperial.ac.uk

N. Bempedelis
School of Engineering and Materials Science
Queen Mary University of London
London E1 4NS, UK
n.bempedelis@qmul.ac.uk

R. Palacios
Department of Aeronautics
Imperial College London
London SW7 2AZ, UK
r.palacios@imperial.ac.uk

S. Laizet
Department of Aeronautics
Imperial College London
London SW7 2AZ, UK
s.laizet@imperial.ac.uk

ABSTRACT

We discuss the flow features from the results of two Bayesian Optimisations (BO) that employ Large Eddy Simulations (LES) as a function evaluation to determine the configuration of two turbines around a 2D hill that yields maximal total power output. Within the LES framework, the turbulent flow field is modelled using the Navier-Stokes equations, while terrain features are captured using an immersed boundary method combined with a wall stress model to account for the viscous sub-layer of the turbulent boundary layer. The turbine wakes are represented by a momentum sink using the actuator disc method. The design variables to optimise in the first case (BO1) are the streamwise location and hub heights of the turbines. BO1 achieves considerable enhancement in power over a reference case by exploiting the local flow acceleration between the hill's speed-up effect and a tall upstream wind turbine. In the second case (BO2) the upstream turbine's hub height is fixed to a more modest value and the tilt angle is added as a design variable to study the benefits of wake steering. A similar positioning is proposed by the optimiser but with 8.72° of disc tilt, meaning a misalignment of the upstream turbine to the inflow reduces its power generation for an overall improvement in the total power of both turbines.

INTRODUCTION

In hilly areas, the turbulent atmospheric boundary layer is significantly affected by terrain topography. In such environments, the flow may exhibit accelerations due to favourable pressure gradients, directional changes, high shear, and increased turbulence intensity resulting from flow separation. An intricate understanding of the flow behaviour is crucial for optimising wind farm performance in complex terrain (Porté-Agel *et al.*, 2020). Several studies have investigated this problem experimentally and with simulations using simple shapes such as the Gaussian and cosine squared hills, which are often used as reference models for real-world topography (Howard *et al.*, 2013; Tian *et al.*, 2013; Yang *et al.*, 2015; Shamsoddin & Porté-Agel, 2017). Zhang *et al.* (2022) examined individual wind turbine wakes over 2D hills of varying slopes, with and without flow separation. They suggest placing the turbine at the hilltop as it was deemed ideal for harvesting more energy and reducing the turbine dynamic loading owing to reduced levels of turbulence. Conversely, turbines placed in the wake of steep hills experienced a decrease in performance due to reverse flow and high levels of turbulence intensity. Liu & Stevens (2020) investigated the trade-offs of wind farm power when leaving a gap in front or behind a 2D hill and concluded that there is no significant effect on the power production, with the hill's recirculation zone being the overriding flow feature. In this study, optimisations are undertaken to enhance the wind farm layout around a 2D hill, considering two turbines: one on top of the hill and one behind, to maximise power production.

The accurate determination of the power output of wind farms in complex terrains relies on the precise prediction of nonlinear and unsteady flow dynamics, covering both atmosphere-to-wake and wake-to-wake interactions. Turbulence-resolving simulations, such as Large Eddy Simulations, effectively capture the pressure gradients, variations in wake trajectory, and flow separation by accounting for the dynamic interactions among the atmospheric boundary layer, complex terrain, and turbine wakes. Recent work by the authors (Jané-Ippel *et al.*, 2023) has validated the use of LES with the open-source flow solver Xcompact3d to model turbulent flow around a constant-section hill, with and without positioning wind turbines around the hill. This validated methodology, capable of representing a broad spectrum of fluid flow phenomena, underpins the optimisations carried out in the present study. The high computational demands of LES limit its use in optimisation tasks, but the integration of Bayesian Optimisation (BO) mitigates this challenge. BO employs surrogate models to approximate the search space, enabling efficient determination of the most informative points for evaluation with a minimal number of function evaluations. This strategic approach facilitates a computationally efficient optimisation process, leveraging the strengths of LES in complex

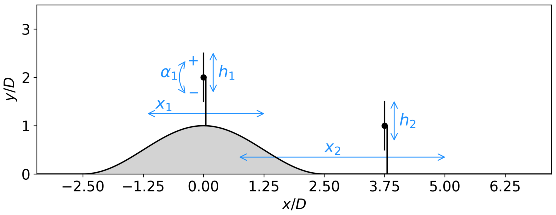

Figure 1. Schematic of the 2D hill with the design variables of BO1: x_1, h_1, x_2, h_2; and BO2: x_1, α_1, x_2, h_2.

terrain wind farm analysis.

This work presents an analysis of flow dynamics and wake-to-wake interactions, utilising time-averaged velocity and turbulence intensity fields from the two optimal configurations identified by the Bayesian Optimisation. We first compare the mean flow in the mid-span plane across three scenarios: a Reference case and the two optimised configurations, BO1 and BO2, highlighting the benefits of strategic turbine positioning in complex terrain. Subsequently, we delve into the turbine blockage in BO1 and wake steering effects in BO2, demonstrating how the upstream turbine's positioning can significantly influence the downstream turbine performance. Lastly, we address the alterations in hill recirculation due to turbine placement, illustrating the nuanced changes in the flow behaviour around the hill and the potential of LES to capture these complex interactions.

METHODOLOGY

Problem Set-up

The complex terrain considered in this work is a constant section hill studied experimentally by Cao & Tamura (2006) defined by a cosine-squared function with a maximum height of $h = 0.04$ m and half-width of $L = 0.1$ m, as depicted in Figure 1 with normalised dimensions. The incoming velocity profile of the experiments is defined by a turbulent boundary layer of height $\delta = 0.25$ m, a friction velocity $u_* = 0.1926$ m/s and a roughness length $y_0 = 0.004$ mm. The actuator disc that represents the wind turbines has a diameter of $D = h$ and a modified thrust coefficient of $C_T' = 1.33$.

The two optimisations performed, BO1 and BO2, have four design variables each. For BO1, the design variables, normalised by the turbine diameter, include streamwise locations ($x_1 \in [-1.25, 1.25]$, $x_2 \in [0.75, 5.00]$) and hub heights ($h_1, h_2 \in [0.70, 1.50]$). BO2 builds upon BO1, fixing the maximum hub height of the upstream wind turbine to $h_1 = 1.0$ and adding tilt as a design variable to compensate for the height restriction ($\alpha_1 \in [-20.0°, 20.0°]$) and explore the merits of wake steering. A constraint is used in both optimisations to ensure a safe turbine operation distance ($x_2 - x_1 > 1.05$). Figure 1 shows a representative sketch of the complex terrain and the design variables of the optimisations for $x_1 = 0.0$, $h_1 = 1.0$, $\alpha_1 = 0°$, $x_2 = 3.75$ and $h_2 = 1.0$. The objective function of the optimisation is the total power, defined as the sum of the average power of each wind turbine. The power of the turbines are normalised by the power generated by a single wind turbine under the influence of the same inlet conditions without the hill's presence.

Computational Fluid Dynamics

The wind farm simulator WInc3D (Deskos *et al.*, 2020), part of the high-order finite-difference framework Xcompact3d (Bartholomew *et al.*, 2020), is used to perform LES of the flow over complex terrain for high Reynolds numbers. The governing equations are the unsteady, incompressible, filtered Navier-Stokes (N-S) equations, based on an explicit LES formulation,

$$\frac{\partial \widetilde{u}_i}{\partial t} + \frac{1}{2}\left(\widetilde{u}_j \frac{\partial \widetilde{u}_i}{\partial x_j} + \frac{\partial \widetilde{u}_i \widetilde{u}_j}{\partial x_j}\right) = -\frac{1}{\rho}\frac{\partial \widetilde{p}}{\partial x_i} - \frac{\partial \tau_{ij}}{\partial x_j} + \nu \frac{\partial^2 \widetilde{u}_i}{\partial x_j \partial x_j} + \frac{F_i}{\rho}, \quad (1)$$

$$\frac{\partial \widetilde{u}_i}{\partial x_i} = 0, \quad (2)$$

where $\widetilde{u}_i = (u_x, u_y, u_z)$ and $\widetilde{p}$ are the spatially filtered components of velocity and pressure fields, $\rho = 1.2\,\mathrm{kg/m^3}$ is the fluid density, τ_{ij} the subfilter-scale stresses and $\nu = 1.5 \times 10^{-5}\,\mathrm{m^2/s}$ the kinematic viscosity. F_i is the body forcing used to model the complex terrain with an Immersed Boundary Method (IBM), and also to model the wind turbines. The subfilter stresses are computed using the standard Smagorinsky model (Smagorinsky, 1963) applying the wall-damping function of Mason & Thomson (1992) to the Smagorinsky constant to avoid excessive dissipation at the surface (Calaf *et al.*, 2010). The discretisation of the governing equations is done using sixth-order compact finite-difference schemes for the spatial derivatives and an explicit third-order Adams-Bashforth for time marching. The use of higher-order schemes necessitates special treatment of the non-linear term of the momentum equation, which is computed in the skew-symmetric form for increased stability and to reduce aliasing errors. The Poisson equation, which guarantees the incompressibility of the velocity field, is fully solved in spectral space via the use of relevant three-dimensional fast Fourier transforms (FFTs). By using the concept of modified wavenumber, the divergence free condition is ensured up to machine accuracy. The pressure mesh is staggered from the velocity mesh by half a node to avoid spurious pressure oscillations observed in a fully collocated approach. More details of the code implementation can be found in Laizet & Lamballais (2009). The terrain features are reproduced with an IBM, which can be combined with a stress wall model to avoid the prohibitively expensive resolution of the viscous sub-layer (Jané-Ippel *et al.*, 2024).

The computational domain has a size of $L_x \times L_y \times L_z = (5 \times 1 \times 1)\delta$. The simulations are run for a total non-dimensional time of $\hat{T} = 20$ to fully converge the second-order statistics after discarding $\hat{T} = 2$ of initial transient. The non-dimensional time unit ($\hat{T}$) is defined as $\hat{T} = Tu_*/\delta$, where T is time. The grid resolution resulting from the validation presented in Jané-Ippel *et al.* (2024) is $n_x \times n_y \times n_z = 385 \times 193 \times 128$, which results in around 9.5M grid nodes. The time step used is $\mathscr{D}\hat{T} \approx 3.85 \times 10^{-5}$. A precursor simulation is performed to generate the neutrally stable turbulent boundary layer that is then used as an inlet boundary condition. A slip wall boundary condition is enforced on the top boundary, while a no-slip wall condition with a wall stress model is used on the bottom boundary. Periodic boundary conditions are applied in the spanwise planes. The outlet boundary condition is applied with one-dimensional convection equations. The convective velocity is calculated as the time and grid-scaled average of the maximum and minimum streamwise velocities at the last interior grid point. The actuator disc method is used to model the wind turbines with a forcing term added to the Navier-Stokes equations and to estimate their power (Jané-Ippel *et al.*, 2024).

The simplicity of the mesh allows an easy implementation of a 2D domain decomposition based on pencils. The computational domain is split into a number of subdomains (pencils), which are each assigned to a message passing interface (MPI) process. The derivatives and interpolations in the x-direction (y-direction, z-direction) are performed in X-pencils, Y-pencils and Z-pencils, respectively. The three-dimensional FFTs required by the Poisson solver are also broken down as series of one-dimensional FFTs computed in one direction at a time. Global transpositions to switch from one pencil to another are performed with the MPI command `MPI_ALLTOALL(V)`. Winc3D can scale well with up to hundreds of thousands of MPI processes for simulations with several billion mesh nodes (Deskos *et al.*, 2020).

Optimisation Algorithm

To optimise the power output of the two-turbine set-up, we adopt a Bayesian Optimisation strategy that utilises Gaussian Process (GP) models as surrogate models, due to its effectiveness with computationally intensive objective function evaluations, such as LES. The optimisation was implemented using the GPyOpt library (The GPyOpt authors, 2016). The method aims to find the optimal set of design variables to maximise the objective function over the design space.

We initiate the optimisation process with a Design of Experiments (DOE), generating initial design points via Latin Hypercube Sampling (LHS) for a comprehensive exploration of the design space. These points are then assessed using high-fidelity simulations to create the first GP model. The GP model, characterised by its mean function and covariance function, is conditioned on training data to predict the mean and uncertainty of the objective function at new points. We utilise the Matérn 5/2 kernel for its proven effectiveness in capturing the dynamics of physical systems (Diessner *et al.*, 2022), with each design variable in the model assigned a unique lengthscale to reflect its specific impact accurately.

Subsequent evaluations are guided by the Lower Confidence Bound (LCB) acquisition function, optimising the balance between exploring new areas and exploiting known regions of the design space. A batch of four design points for simultaneous evaluation is selected to incorporate new data and update the GP model. This approach ensures efficient progression towards finding the optimal design variables within the computational constraints of the evaluations.

Each evaluation based on one LES utilised 2048 CPU cores from the ARCHER2 UK National Supercomputing Service, equipped with dual AMD EPYC^{TM} 7742 64-core processors. The wall clock time per each LES was approximately 3.6 hours. Each optimisation performed 100 evaluations which result in a cumulative computational cost around 737,280 CPUh. Each optimisation required approximately five days of wall clock time to complete, accounting for parallel initial sampling and the use of batches of four evaluations at each surrogate model update. Convergence in both BO1 and BO2 was assessed by monitoring the cumulative number of simulations until the improvement in the maximum power output dropped below a threshold of 1%. Under this criterion, BO1 converged after 56 simulations and BO2 after 40 simulations. The computational demands and time associated with the BO algorithm were marginal compared to the LES evaluations.

RESULTS

In this section, we present comparisons among three configurations: two derived from Bayesian Optimisations (BO1 and BO2) and a Reference case. The Reference case, a standard arrangement, situates one wind turbine at the centre of a hill with another turbine 5D downstream. In this setup, the normalised total power output is $P_{Rc} = 2.70$, each unit reflecting the power generated by a single turbine under the same inflow conditions on flat terrain.

For BO1, the optimal design variables found ($x_1 = 0.12$, $h_1 = 1.5$, $x_2 = 1.17$, $h_2 = 1.02$) position the upstream turbine slightly to the leeward side of the hill's centre, elevated at a higher hub height. The downstream turbine is located 1.05D away from the upstream one, satisfying the optimisation constraint of the distance between the turbines. This layout facilitates a total power output of $P_{BO1} = 4.66$, marking a substantial 73% enhancement over the Reference case. BO2 aimed to mimic similar conditions through modifications in the upstream turbine's tilt angle rather than hub height, thus exploring wake steering effects. The optimal design variables found for BO2 ($x_1 = -0.12$, $\alpha_1 = 8.72°$, $x_2 = 1.39$, $h_2 = 0.98$) locate the upstream turbine slightly ahead of the hill's centre, with a tilt angle misaligning the disc with the inflow. The downstream turbine, positioned 1.51D away, allows more space for the upstream wake to be deflected, improving its conditions. This configuration yields a power output of $P_{BO2} = 3.61$, demonstrating a 34% increase over the Reference case.

Mean flow in the midspan-plane

To gain insight into the differential performance across scenarios, we assess the mean velocity (Figure 2) and turbulence intensity (Figure 3) fields in the plane crossing the aligned turbine centres.

Figure 2 presents the normalised time-averaged streamwise velocity field along the mid-plane. Black dashed lines demarcate the recirculation boundaries, identified by zero-velocity contours that define the edge of reverse flow zones. In the reference configuration, the bulk of the total power generated comes primarily from the upstream turbine which exploits the flow speed-up due to the favourable pressure gradient on the hill's front side, contributing $P_{Rc_U} = 2.33$, while the downstream turbine adds a minor $P_{Rc_D} = 0.37$. As depicted in Figure 2a, the downstream turbine suffers from both the low velocity induced by the hill and the upstream turbine's wake, resulting in lower power output. Conversely, as illustrated in Figures 2b and 2c, both optimal scenarios locate the downstream turbine in positions clear of the hill wake. In these layouts, the downstream turbine exploits the local flow acceleration between the hill speed-up and the upstream turbine wake. The placement of the downstream turbines closer to the hill's flow separation point, modifies the hill's recirculation pattern in the mid-span plane.

In the BO1 configuration, adjustments in the upstream hub height and in the turbine placement significantly enhance the downstream turbine's output to $P_{BO1_D} = 2.35$, marginally surpassing the upstream turbine's $P_{BO1_U} = 2.31$. This suggests that the upstream turbine's blockage partially redirects the flow towards the downstream turbine. In the BO2 setup, the upstream turbine, misaligned with the inflow, generates less power $P_{BO1_U} = 2.14$ compared to the reference and BO1 scenarios where the turbines align with the inflow. However, in BO2, the downstream turbine, overlapping with the upstream wake due to hub height limitations, still manages to produce $P_{BO2_D} = 1.47$, benefiting from the modified flow dynamics.

Additional understanding can be derived from the analysis of second-order statistics. Figure 3 illustrates the normalised time-averaged streamwise turbulence intensity across the mid-plane in each scenario. In the Reference scenario, the

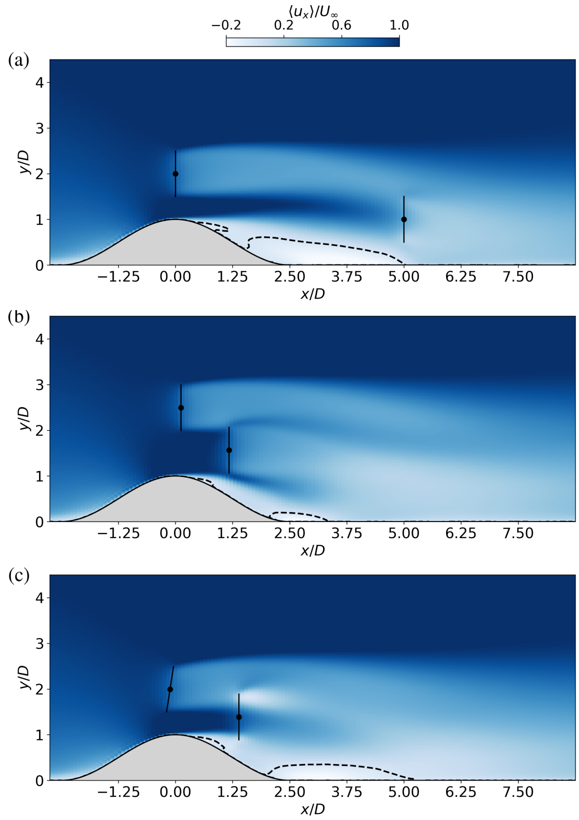

Figure 2. Normalised time-averaged streamwise velocity field of the Reference case (a), BO1 (b) and BO2 (c). The zero velocity contours are shown as black dashed lines.

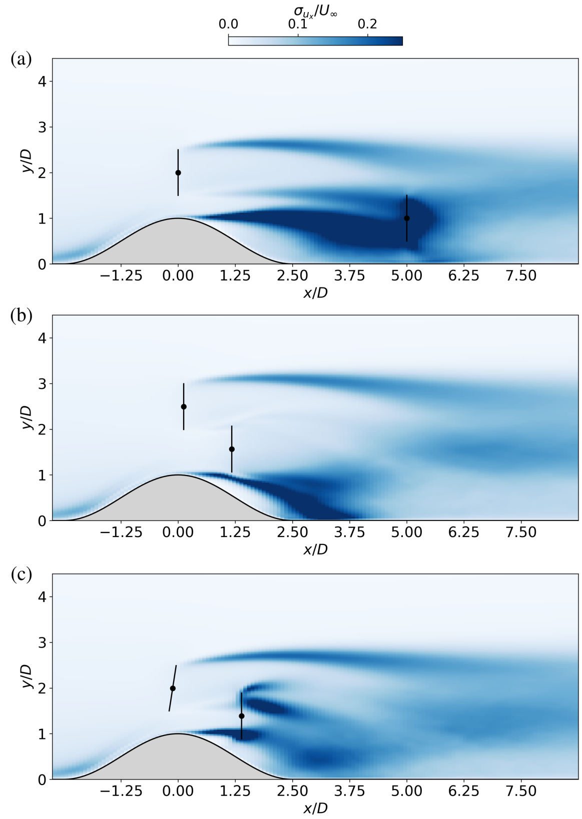

Figure 3. Normalised time-averaged streamwise turbulence intensity field of the Reference case (a), BO1 (b) and BO2 (c).

downstream turbine suffers from substantial turbulence due to flow separation at the hill combined with turbulence that emerges from the edges of the upstream turbine. Conversely, the turbines in BO1 benefit from comparatively undisturbed conditions. The BO2 setup shows the impact from the wakes of both the hill and the upstream turbine, although less than in the Reference scenario. Notably, the turbulence intensity increases when the wakes of the upstream and downstream turbines intersect in the BO2 configuration. In both the Reference and BO2 scenarios, such conditions could adversely affect the downstream turbine through increased dynamic loadings due to turbulence. However, the turbines in optimal configurations demonstrate a reduction in the turbulent kinetic energy associated with the hill wake relative to the Reference case.

Turbine Blockage in BO1

The increased power output of the downstream turbine in the BO1 setup can be understood by its strategic positioning, which avoids the wakes of both the upstream turbine and the hill. However, the optimisation results indicate that it is placed near the upstream turbine, prompting a closer look at the blockage effect of the upstream turbine. We compare the flow fields from simulations with and without the upstream turbine to analyse this influence. The presence of the upstream turbine leads to a 5% increase in the power of the downstream turbine, a result of the flow being redirected by the upstream turbine's blockage.

Figure 4 illustrates the velocity deficit $\langle u_{\mathrm{def}} \rangle$ evaluated in a streamwise plane 0.25D upstream of the downstream turbine ($x/D = 0.92$). The velocity deficit is defined as $\langle u_{\mathrm{def}} \rangle = \langle u_{\mathrm{nt}} \rangle - \langle u_{\mathrm{t}} \rangle$, where $\langle u_{\mathrm{nt}} \rangle$ is the time-averaged streamwise velocity field in the absence of the upstream turbine and $\langle u_{\mathrm{t}} \rangle$ is the mean streamwise velocity with the upstream turbine, both normalised by U_∞ (U_∞ is the time-averaged streamwise velocity of the turbulent boundary layer at the top boundary). The horizontal line at $y/D = 0.7$ marks the hill's edge at this longitudinal location, and the dashed lines show the contours of velocity deficit at -0.02 and -0.04. These contours indicate regions where the flow speed is increased by 2% and 4% of U_∞ respectively. The turbines' centres and the extent of their rotor discs are projected onto the plane for scale. The influence of the upstream turbine is visible as a positive velocity deficit, slightly covering the downstream turbine's projected area. The contours reveal areas of faster flow close to the upstream turbine and along the hill edge. This flow acceleration spreads across a large area impacting the downstream turbine, which explains the observed performance boost. The acceleration is particularly strong near the hill, indicating that the upstream turbine's presence has altered the hill's recirculation zone.

Wake Steering in BO2

Wake steering, a method to control the aerodynamic interaction between turbines, is exploited in BO2 by introducing a tilt angle to the upstream turbine. This intentional deflection of the wake can mitigate detrimental effects on downstream turbines when there is wake overlapping, improving the overall efficiency of wind farms. In BO2, the upstream turbine has a tilt of $\alpha_1 = 8.72^\circ$, which, despite reducing its power output by 6% if compared to the no-tilt case, leads to a substantial increase of 23.5% for the downstream turbine's power. The collective effect of this wake steering approach is a net gain of 4.3% in total power compared to the no-tilt configuration, a significant improvement considering the close proximity of

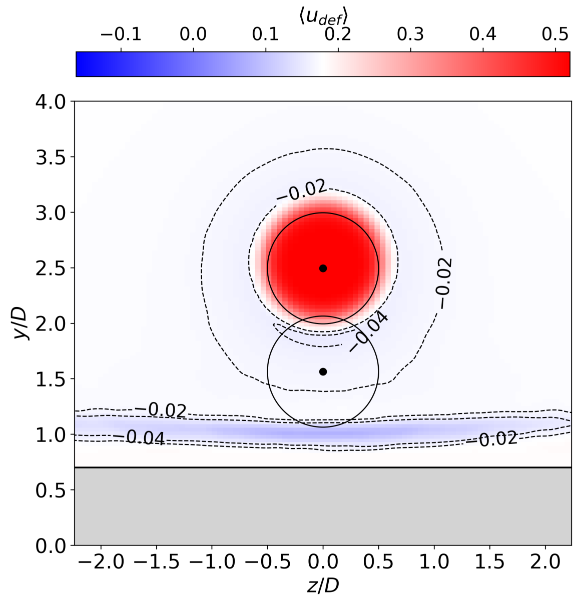

Figure 4. Velocity deficit caused by the upstream turbine in BO1 evaluated 0.25D upstream of the downstream turbine.

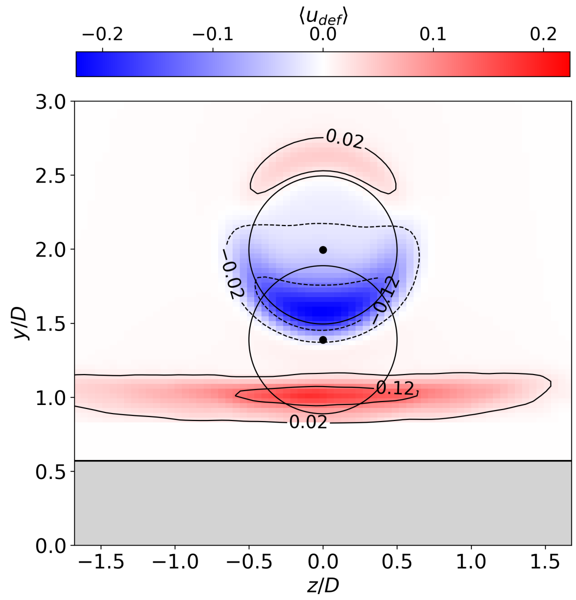

Figure 5. Velocity deficit in the wake of the upstream turbine induced by its tilt in BO2 evaluated 0.25D upstream of the downstream turbine.

the turbines.

Figure 5 illustrates the alterations in the flow pattern due to the upstream turbine's tilt, applying the velocity deficit concept as defined in the previous section. The velocity deficit $\langle u_{\mathrm{def}} \rangle$ is computed by contrasting the flow with and without the upstream turbine's tilt at a plane 0.25D upstream of the downstream turbine ($x/D = 1.139$). The horizontal line at $y/D = 0.57$ designates the hill's edge, and the dashed lines delineate the contours of velocity deficit at -0.12, -0.02, 0.02, and 0.12. Tilting diverts the wake upwards, creating an accelerated flow in the lower region of the wake's trajectory, which benefits the downstream turbine over a considerable area. Concurrently, as also seen in Figure 4, the tilt induces modifications in the hill's recirculation, notably a horizontal deceleration over the hill edge, which has a marginal adverse effect on the downstream turbine's performance. The interaction of these factors underscores the utility of wake steering in enhancing overall energy capture in complex terrain settings.

Hill Recirculation Changes

The wind tunnel experiments by Cao & Tamura (2006) denote the hill's natural separation point occurring shortly after the crest in the absence of turbines. Figure 2 illustrates how turbine placement substantially influences the flow patterns around the hill in the mid-span plane. The interaction with turbine wakes not only accelerates the flow but also alters the hill's recirculation characteristics when turbines are placed near the separation point. Figure 6 provides a 3D representation of the recirculation edge, delineated by the zero-velocity contour and juxtaposed with the actuator discs for each case.

In the Reference case (Figures 2a and 6a), the upstream turbine's positioning on the hill's summit causes the flow to reattach in a small area of the leeward side of the hill. The introduction of a downstream turbine, as seen in BO1 and BO2, enlarges the reattachment zone owing to the blockage effects exerted by both turbines. However, as evident in Figure 3, changes in the recirculation pattern arise not solely due to the downstream turbine but are also a consequence of the upstream turbine's blockage. It is noteworthy that despite the similarities in turbine placement between BO1 and BO2, the extent of the recirculation edge near the mid-span region varies significantly.

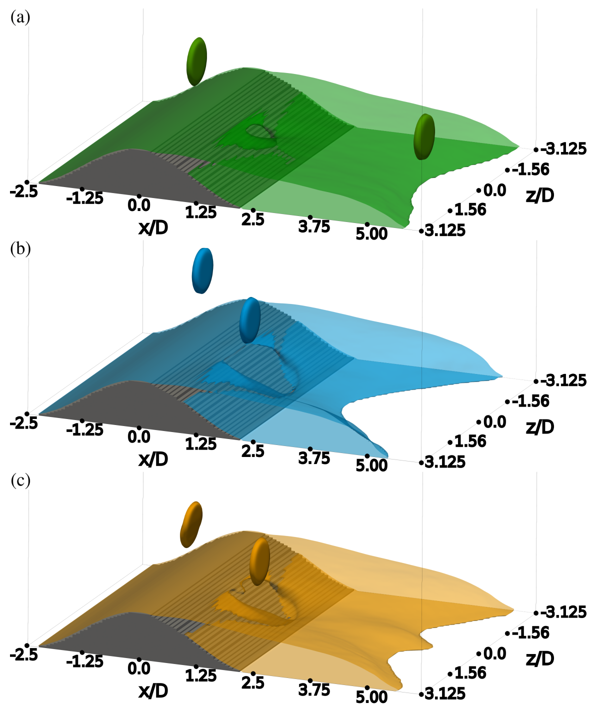

Figure 6. 3D visualisation of the hill recirculation edge depicted by the contour of zero-velocity for the Reference (a), BO1 (b), and BO2 (c) cases.

This complex interaction is critical to understand as it is often not captured by lower-fidelity models such as the superposition method, leading to inaccuracies in power prediction for turbines influenced by flow separation (Liu & Stevens, 2020). LES becomes indispensable in these scenarios, capturing non-linear effects and providing the necessary fidelity for the optimisation of turbine positioning within such complex terrains with flow separation.

CONCLUSIONS

This study has established Bayesian Optimisation based on Large Eddy Simulations as an efficient approach for improving the power output of wind tubrines located in complex terrains. We have demonstrated the significant benefits of optimising turbine placement, particularly in the typically underused area behind the hill, by harnessing the accelerated flow caused by the hill's shape. As seen in the two optimal cases, BO1 and BO2, the turbines' strategic locations reduce wake interference, maximising the capture of energy from the local accelerated flow. Adjustments such as increasing the upstream turbine's tower height (BO1) or employing tilt angles upstream (BO2) have resulted in marked power gains with respect to a Reference case. Furthermore, our findings show the impact of turbine placement on hill recirculation and the subsequent aerodynamic interactions within wind farms in complex terrain. The use of LES has been pivotal, as it can model the complex non-linear phenomena that elude lower-fidelity approaches. This advantage is essential for the precise prediction of power generation and the optimisation of turbines in environments where wake interactions dominate.

Future work will expand the application of this methodology to study a wind farm in realistic complex terrains, involving a greater number of turbines than the current two-turbine in a simplified hill model. The high-fidelity approach will be crucial for investigating the effects of wake steering in scenarios where wake overlaps are prominent. This progression will allow for a more comprehensive understanding of aerodynamic interactions and their optimisation in real-world wind farm configurations.

ACKNOWLEDGEMENTS

This work used the ARCHER2 UK National Supercomputing Service via a Pioneer project and resources provided by the UK Turbulence Consortium (EP/R029326/1). The first author thanks NERC and SSE Renewables Services (UK) Limited for funding a PhD studentship in the Department of Aeronautics at Imperial College London via the Science and Solutions for a Changing Planet DTP (NE/S007415/1).

REFERENCES

Bartholomew, P., Deskos, G., Frantz, R. A. S., Schuch, F. N., Lamballais, E. & Laizet, S. 2020 Xcompact3d: An open-source framework for solving turbulence problems on a cartesian mesh. *SoftwareX* **12**, 100550.

Calaf, M., Meneveau, C. & Meyers, J. 2010 Large eddy simulation study of fully developed wind-turbine array boundary layers. *Physics of Fluids* **22**, 015110.

Cao, S. & Tamura, T. 2006 Experimental study on roughness effects on turbulent boundary layer flow over a two-dimensional steep hill. *Journal of Wind Engineering and Industrial Aerodynamics* **94**, 1–19.

Deskos, G., Laizet, S. & Palacios, R. 2020 Winc3d: A novel framework for turbulence-resolving simulations of wind farm wake interactions. *Wind Energy* **23**, 779–94.

Diessner, M., O'Connor, J., Wynn, A., Laizet, S., Guan, Y., Wilson, K. & Whalley, R. D. 2022 Investigating bayesian optimization for expensive-to-evaluate black box functions: Application in fluid dynamics. *Frontiers in Applied Mathematics and Statistics* **8**.

Howard, K. B., Chamorro, L. P. & Guala, M. 2013 An experimental case study of complex topographic and atmospheric influences on wind turbine performance. *51st AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition 2013* .

Jané-Ippel, C., Bempedelis, N., Palacios, R. & Laizet, S. 2023 High-fidelity simulations of wake-to-wake interaction in an atmospheric boundary layer over a complex terrain. *Journal of Physics: Conference Series* **2505** (1), 012033.

Jané-Ippel, C., Bempedelis, N., Palacios, R. & Laizet, S. 2024 Bayesian optimisation of a two-turbine configuration around a 2d hill using large eddy simulations. *submitted to Wind Energy* .

Laizet, S. & Lamballais, E. 2009 High-order compact schemes for incompressible flows: A simple and efficient method with quasi-spectral accuracy. *Journal of Computational Physics* **228**, 5989–6015.

Liu, L. & Stevens, R. J. A. M. 2020 Effects of two-dimensional steep hills on the performance of wind turbines and wind farms. *Boundary-Layer Meteorology* **176**, 251–69.

Mason, P. J. & Thomson, D. J. 1992 Stochastic backscatter in large-eddy simulations of boundary layers. *Journal of Fluid Mechanics* **242**, 51–78.

Porté-Agel, F., Bastankhah, M. & Shamsoddin, S. 2020 Wind-turbine and wind-farm flows: A review. *Boundary-Layer Meteorology* **174**, 1–59.

Shamsoddin, S. & Porté-Agel, F. 2017 Large-eddy simulation of atmospheric boundary-layer flow through a wind farm sited on topography. *Boundary-Layer Meteorology* **163**, 1–17.

Smagorinsky, J. 1963 General circulation experiments with the primitive equations: I. the basic experiment. *Monthly Weather Review* **91** (3), 99 – 164.

The GPyOpt authors 2016 GPyOpt: A Bayesian Optimization framework in Python. http://github.com/SheffieldML/GPyOpt.

Tian, W., Ozbay, A., Yuan, W., Sarakar, P. & Hu, H. 2013 An experimental study on the performances of wind turbines over complex terrain. *51st AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition 2013* .

Yang, X., Howard, K. B., Guala, M. & Sotiropoulos, F. 2015 Effects of a three-dimensional hill on the wake characteristics of a model wind turbine. *Physics of Fluids* **27**, 025103.

Zhang, Z., Huang, P., Bitsuamlak, G. & Cao, S. 2022 Large-eddy simulation of wind-turbine wakes over two-dimensional hills. *Physics of Fluids* **34**, 065123.